

AD-757 117

FIRST QUARTERLY TECHNICAL REPORT,
QUARTER ENDED 31 JANUARY 1973

Clinton W. Kelly, III, et al

Decisions and Designs, Incorporated

Prepared for:

Advanced Research Projects Agency
Office of Naval Research

March 1973

DISTRIBUTED BY:

NTIS

National Technical Information Service
U. S. DEPARTMENT OF COMMERCE
5285 Port Royal Road, Springfield Va. 22151


AD757117

First Quarterly Technical Report

**Quarter ended 31 January 1973
Contract N00014-73-C-0149**


**Sponsored by
Advanced Research Projects Agency
ARPA Order No. 2271 dated 7 August 1972**

Reproduced by
**NATIONAL TECHNICAL
INFORMATION SERVICE**
U S Department of Commerce
Springfield VA 22151



Decisions and Designs, Incorporated

Suite 600, 7900 Westpark Drive
McLean, Virginia 22101
(703) 821-2828



DOCUMENT CONTROL DATA - R & D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) Decisions and Designs, Inc. Suite 600, 7900 Westpark Drive McLean, Virginia 22101		2a. REPORT SECURITY CLASSIFICATION Unclassified	
		2b. GROUP ---	
3. REPORT TITLE First Quarterly Technical Report Quarter Ended 31 January 1973 Contract N00014-73-C-0149			
4. DESCRIPTIVE NOTES (Type of report and inclusive dates) Quarterly 1 October 1972 through 31 January 1973			
5. AUTHOR(S) (First name, middle initial, last name) Dr. Clinton W. Kelly, III Dr. Cameron R. Peterson			
6. REPORT DATE Submitted March 1973	7a. TOTAL NO. OF PAGES 50	7b. NO. OF REFS One	
8a. CONTRACT OR GRANT NO. N00014-73-C-0149	9a. ORIGINATOR'S REPORT NUMBER(S) 8-1		
b. PROJECT NO. ARPA Order No. 2271 dated 7 Aug 1972			
c.	9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report) --		
d.			
10. DISTRIBUTION STATEMENT Distribution of this document is unlimited.			
11. SUPPLEMENTARY NOTES -----		12. SPONSORING MILITARY ACTIVITY Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209	
13. ABSTRACT <p>Task 1 — Investigate procedures for improving human judgments of probabilities & utilities for decision making. In performing this investigation, DDI will conduct research on the application of decision theory to problems in policy analysis & resource allocation. This includes on-line, case study oriented research with decision makers for the purpose (a) of determining strengths & weaknesses in present decision theoretic technology, & (b) for promoting the use of decision theoretic concepts through the familiarization of decision makers with these concepts. This task also includes laboratory research on procedures for improving human judgments on probabilities & utilities. That is, research on procedures for (a) encoding uncertainties as probabilities & (b) incorporating attitudes toward risk into utilities.</p> <p>Task 2 — Conduct problem oriented workshops for DOD personnel in which the potential value of decision analysis techniques is displayed to decision makers by showing them how decision analysis can be applied to real problems.</p> <p>Task 3 — Prepare a handbook for users of decision analysis designed for the manager, or staff, responsible for organizing and managing a decision analysis rather than for the decision analytic technician.</p>			

Ia

Tb

First Quarterly Technical Report

Quarter ended 31 January 1973
Contract N00014-73-C-0149

Sponsored by
Advanced Research Projects Agency
ARPA Order No. 2271 dated 7 August 1972

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency or the U. S. Government.

Decisions and Designs, Inc.
Suite 600, 7900 Westpark Drive
McLean, Virginia 22101

DATA PAGE

This research was supported by the Advanced
Research Projects Agency of the Department
of Defense and was monitored by ONR under
Contract No. N00014-73-C-0149

Identifying Data:

ARPA Order No.	: 2271 dated 7 August 1972
Program Code No.	: 3D20
Name of Contractor	: Dccisions and Designs, Inc.
Effective Date of Contract	: 1 October 1972
Contract Expiration Date	: 31 December 1973
Amount of Contract	: \$123,455
Contract Number	: N00014-73-C-0149
Principal Investigators and Telephone Number	: Dr. Clinton W. Kelly, III and Dr. Cameron R. Peterson (703) 821-2828
Scientific Officer	: Dr. Martin A. Tolcott Director, Engineering Psychology Projects Office of Naval Research
Short Title of Work	: Decision Theory Research

Table of Contents

	<u>Page</u>
1. SUMMARY	1-14
1.1 Introduction	1-3
1.1.1 Technical problem	1-2
1.1.2 General methodology	2-3
1.2 Technical Results	3-14
1.2.1 Task 1 - research on the application of decision theory to policy analysis	3-8
1.2.2 Task 1 - research on the application of decision theory to resource allocation	8-11
1.2.3 Task 1 - technical support experiment	11-12
1.2.4 Task 2- workshop	12
1.2.5 Task 3 - handbook	13-14
2. TASK 1 - RESEARCH ON THE APPLICATION OF DECISION THEORY	15-40
2.1 Action Selection (Primary Decisions)	15-22
2.1.1 Case study 1 - decision analysis with probabilities and utilities	15-20
2.1.2 Case study 2 - decision analysis with probabilities only	20-22
2.2 Information Seeking Decisions (Secondary Decisions)	22-33
2.2.1 Difficulties in using the primary-decision approach	23-24
2.2.2 An approach based on direct assessment of value - use of value diagrams	24-29
2.2.3 Value diagram - first case study	29-31

Contents (continued)	<u>Page</u>
2.2.4 Value as a function of decision variables - second case study	31-33
2.3 Supporting Technical Experiments	33-40
2.3.1 Experiment 1 - odds versus probabilities as the response mode for assessing the likelihood of categorical events	33-36
2.3.2 Experiment 2 - probabilities of points versus intervals along a continuum	36-40
2.3.3 Experiment 3 - decomposition in the assessment of continuous variables	40
3. TASK 2 - WORKSHOPS	41
4. TASK 3 - HANDBOOK FOR USERS OF DECISION ANALYSIS	42-45

First Quarterly Technical Report
Contract N00014-73-C-0149

Decision Theory Research

1. SUMMARY

1.1 Introduction

This is the first quarterly technical report submitted under ONR contract N00014-73-C-0149. It describes research conducted from the inception of the contract through 31 January 1973, as set forth in the above named contract and as amplified by Decisions and Designs, Inc. in Proposals 10-72 of 26 July 1972 and 12-72 of 22 August 1972.

1.1.1 Technical Problem - Three general research tasks were contracted for.

Task 1. Investigate procedures for improving human judgments of probabilities and utilities for decision making. In performing this investigation, DDI will conduct research on the application of decision theory to problems in policy analysis and resource allocation. This includes on-line, case study oriented research with decision makers for the purpose (a) of determining strengths and weaknesses in present decision theoretic technology, and (b) for promoting the use of decision theoretic concepts through the familiarization of decision makers with these concepts. This task also includes laboratory research on procedures

/

for improving human judgments on probabilities and utilities. That is, research on procedures for (a) encoding uncertainties as probabilities and (b) incorporating attitudes toward risk into utilities.

Task 2. Conduct problem oriented workshops for DOD personnel in which the potential value of decision analysis techniques is displayed to decision makers by showing them how decision analysis can be applied to real problems.

Task 3. Prepare a handbook for users of decision analysis designed for the manager, or staff, responsible for organizing and managing a decision analysis rather than for the decision analytic technician.

1.1.2 General Methodology - The underlying methodology guiding the research effort was the interaction of the investigators with DOD decision makers as they worked on current decision problems. This was done for two reasons: first, to insure the relevance of the research effort, and second, to introduce various high level DOD personnel to decision theoretic concepts and to encourage the use of these concepts in their daily work. Although seminars and other formal instructional techniques can be used to explain the concepts of decision theory to DOD personnel, we believe a far more effective procedure is to demonstrate the utility of those concepts by showing the individual how they can be used to solve his current problems. Thus, in addition to conducting basic research and experimentation, the investigators are functioning as change agents promoting the use of decision theory

by illustrating its applicability to everyday problems of concern to the decision maker.

We believe the researcher's utility will increasingly depend on his capabilities to encourage clients to practice new decision making technologies through joint participation in problem solving, through creating means for self-discovery by the client, and by encouraging, within client organizations, the development of good students who can become centers of innovation.

1.2 Technical Results

1.2.1 Task 1 - Research on the application of decision theory to problems in policy analysis. Initial research conducted with policy analysts showed that they were lacking a conceptual basis for evaluating the impact of policy alternatives on desired outcomes or goals. Uncertainties are most frequently incorporated into a policy analysis by treating them as certainties and, to this extent, the analyses fail to provide decision makers with an assessment of the risk associated with policy alternatives.

The intelligence support provided to policy makers is structured in such a way that the policy makers cannot derive maximum benefit from it. Most of the intelligence estimates take the form of a narrative description of one possible state of affairs, a single scenario outlook. Therefore, for the most part, these intelligence analyses do not reflect the impact of any policy alternative which the policy analyst might

be considering. To be of maximum benefit, intelligence should, at least in part, be a standing reflection of the policy maker's capacity to influence events. To the extent that intelligence estimates do not satisfy this function, policy analysts attempt to bridge the gap, either by performing their own intelligence analyses or by trying to modify the existing analyses. The danger inherent in either course of action is that the policy analyst may unconsciously bias the intelligence analysis in order to promote a favored policy alternative. That is, the policy analyst may confuse judgments about the likelihood of events with preferences for the consequences of an event.

As a prescriptive theory of policy analysis, a decision theoretic approach would require: first, that the likelihoods associated with uncertain future events affecting the outcome of a policy be encoded as probabilities; and second, that the decision maker's preference for each possible consequence be encoded as a multi-dimensional utility. In this way, decision theory would first separate judgments about the likelihood of an event occurring from assessments about the desirability of an event, and replace qualitative analysis by quantitative analysis. In both areas, therefore, this would represent a substantial change from current practice.

At the functional level, by distinguishing between judgments about probabilities and those of preferences for consequences, roles are defined for both the intelligence analyst and the policy analyst. The intelligence analyst, or technical expert, should be responsible

for providing the individual probability assessments, and the policy analyst should be concerned with transforming goals to policy alternatives, assessing preferences for the consequences of these alternatives, and recommending a policy for implementation.

At the substantive level, although arguments can be advanced for the adoption of a total decision theoretic approach to policy analysis (one in which the policy analyst uses formal procedures for assessing preferences and performing trade-offs among alternative policies), we believe this goal is not yet practical, both because of current limitations to decision theoretic technology and because of the administrative problems which arise anytime an ongoing process is substantially modified. We do believe, however, that some immediate measures can be taken which are within the spirit of decision theory and which will improve policy analysis. In brief, these measures consist of (a) having policy analysts inform intelligence analysts about alternatives they are evaluating, i.e., the establishment of what Edwards (1973) calls "retrograde information flow" between policy analysts and intelligence analysts, and (b) having the intelligence analysts supply the policy analysts with probability assessments conditional upon these policy alternatives.

Adopting the above measures should rectify what Hughes (1959) called, "a growing crisis of relevance" where the quantity of intelligence provided to the policy analyst, instead of aiding in the selection and evaluation of policy alternatives, may have the opposite result. We believe this

lack of relevance is due to a misinterpretation of the philosophy pervading both the policy and intelligence communities which states that intelligence should remain independent of policy making. In our opinion, this statement implies that intelligence should not confuse the likelihood of an event occurring with its desirability if it occurs. Unfortunately, it is frequently interpreted to mean that the two communities should not communicate. Thus, intelligence analysts write about what they perceive to be important, intelligence remains independent of policy making and policy making is often therefore independent of intelligence.

By providing policy analysts with assessments of the likelihood of uncertain events which could impact upon the outcome of each policy alternative, intelligence becomes more relevant to the policy analyst. First, because probabilities are used to describe degrees of certainty versus words, the individual probability assessments can be combined in a probability model to compute probability distributions over the various dimensions associated with each consequence; second, since each assessment is conditional upon a particular policy alternative, these probability distributions allow the policy analyst to informally select a course of action or policy alternative by balancing off various benefits and costs in consideration of the risk involved. This selection process, although accomplished informally, should constitute an improvement over the current procedures inasmuch as the trade-offs can now be discussed among policy analysts with full knowledge of the multi-dimensional risks associated with each potential course of

action.

To determine the practicality, both technically and administratively, of implementing the above procedures, a case study was performed and is described in detail in Section 2.1.2. The conclusion of this case study is that assessing probabilities for events, conditional upon policy alternatives, is both feasible and desirable. Policy analysts were able to describe policy alternatives to intelligence analysts in an operationally meaningful way. The intelligence analysts, on the other hand, were able to evaluate the impact of policy alternatives on the likelihood of forecast events. This is significant in that such evaluation is a different task from that normally performed by the intelligence analysts. In addition, the intelligence analysts were able to identify leverage points which outlined areas for further policy consideration. For their part, policy analysts could understand the significance of the conditional probability assessments and seemed to be willing to accept them as useful outputs from the intelligence analyst. In general, the results of the case study indicate that creative estimating could go much further than it now does in discerning the possibilities for policies. However, we believe that the intelligence estimators' incentive toward the risk-taking inherent in making such estimates, is such, that considerable encouragement from policy makers will be required in order to get them to participate actively in the process.

Further research in policy analysis will be oriented toward the investigation of a variety of problems which must be solved before

it would be practical for policy analysts to begin formally using decision theory in making policy trade-offs. These problems include decisions as to whose preferences should be used in an analysis, political problems associated with making trade-offs and preferences explicit, and the development of practical procedures for assessing preferences for multi-dimensional outcomes.

1.2.2 Task 1 - Research on the application of decision theory to resource allocation. Research on resource allocation has been directed toward the problem of allocating resources for intelligence collection. Three case studies have been carried out and are described in Sections 2.1.1, 2.2.3, and 2.2.4.

In conducting the resource allocation case studies, two conclusions were reached. First, in making budgeting decisions, the increment, rather than the base, receives most of the review effort. Budget decisions, despite the implementation of the program planning and budgeting system (PPBS), still are made in an arbitrary and exclusionary way. In the area of systems for intelligence collection, security constraints contribute to this problem. Second, techniques for the evaluation of intelligence collection systems have generally been based on measurements of "technical variables". Little emphasis has been placed on "user-oriented variables". For these reasons, no way has been found to effectively plan programs or to judge accomplishments of these programs effectively. While we, as decision theorists, have little more to suggest than anyone else, relative to correcting basic faults in the budget

process, we do believe that decision theory can uniquely furnish the basis for assessing the value of collection systems from a user-oriented perspective.

In theory, the collection of intelligence information should be governed by a principle from decision theory which states,

In principle, it is worthwhile to buy information or- systems to collect that information only if that information may serve to change your behavior, and the value of the information is exactly the difference between the expected values of the old and new behaviors.

Thus, decision theory says that requirements for intelligence should stem from an analysis of potential decisions which may be sensitive to that intelligence. In practice, this principle seems to be ignored. Schlessinger (1972) sums up the present problem in intelligence collection by saying,

The consumer frequently fails to specify his product needs for the producer; the producer, uncertain about eventual demands, encourages the collector to provide data without selectivity or priority; and the collector emphasizes quantity rather than quality.

Many people believe that this excessively conservative approach toward the collection of intelligence has resulted in the expenditure of far more resources for intelligence collection than is really necessary. Most of those responsible for collecting intelligence appear to have adopted what Hammond (1971) has called a, "proxy-conservative" approach to decision making. They have paid far more attention to the value of collecting information about a particular event, given

that the event occurs, than they have about the likelihood of the event occurring or about the impact of the information on a primary decision.

The problem with implementing a decision theoretical information value analysis is that collection decisions often must be made in the absence of a primary decision problem. That is, it is often the case that no primary decision can be identified a priori. For this reason, the first task undertaken was to develop a model for determining the value of information in the absence of a primary decision. This model will be described and illustrated in a forthcoming technical report by applying it to a hypothetical problem concerned with assessing the relative value of several collection systems collecting on the same geographical area. This hypothetical problem will serve to illustrate features of several real problems which were studied during the first quarter. Portions of the above methodology were evaluated in the case studies described in Sections 2.1.1 and 2.2 below.

The intelligence value methodology has two key phases. First, the methodology requires that the decision variables, or control variables, of the allocation problem, be identified. The value of any collection system is then assessed as a function of these decision variables. Second, the methodology requires that intelligence value be assessed by three independent but correlated procedures. These procedures are based on 1) information demands as perceived by the intelligence analysts who utilized the information provided by the collection systems, 2) requirements for finished intelligence, articulated by higher authority,

and 3) by military importance of activities located in specific geographical areas. The output of the intelligence value analysis shows the value of each alternative system as a function of the decision variables which are available to the decision maker for manipulation. Heretofore, most analyses had not been conducted in such a way that the impact of the decision variables was explicit. Just as we believe the assessments of probabilities by intelligence analysts should be displayed as a function of policy alternatives, similarly we believe that the assessment of value of alternative intelligence collection systems should be expressed as a function of those parameters available to the decision maker for control.

We experienced little difficulty in getting decision makers to accept the idea of assigning an intelligence value to alternative collection systems. However, considerable resistance was encountered in attempting to apply a formal procedure for resource allocation to the various problems once the value of the alternative systems had been assessed. As was the case with policy analysis, the decision makers felt far more comfortable working with the trade-offs intuitively, rather than explicitly stating their preferences and using a formal procedure for selecting from among the alternative systems.

1.2.3 Task 1 - Technical support experiment. Subjects were run during this quarter on two experiments designed to evaluate response modes for probability assessment. The first experiment evaluated the relative value of probabilities versus odds as response modes for assessing the likelihood of categorical events. The second

experiment evaluated two different response modes for assessing the relative likelihoods of events that were on a continuum. The task of probability assessment is markedly different for the two classes of events. Categorical variables usually include only a few events that differ qualitatively. For example, a plane may be either a friend or an enemy. A coup may either occur or not occur. Or a bomb may either miss a ship, hit it and cause damage, or hit it and sink it. A continuous variable, on the other hand, consists of quantitative events, such as the top speed of an airplane or the size of a budget, in dollars. The results of these two experiments are now being analyzed and it is expected that they will be written up as technical reports during the next quarter.

1.2.4 Task 2 - Workshop. Participants for a series of workshops on decision theory have been tentatively identified. Our experience to date, however, shows that if these workshops are to be really useful, it is not sufficient for them to consist of a tutorial exposition of decision theory. Therefore, it has been decided that the format for these workshops should be current problem oriented. It is planned that prior to each workshop, a decision problem of one of the participants will be taken and used as a vehicle to illustrate how decision theory can be applied to real problems. The participants for each workshop will be selected so that they will have a knowledge of and interest in, the particular problem which will be studied.

1.2.5 Task 3 - Handbook. A large amount of the effort during the preceding quarter was devoted in developing a case study overview of decision analysis that will serve as the introductory chapter of the handbook. This overview consists of a concrete decision of whether or not a ship's captain should shoot down an approaching plane that may be either friend or foe. The analysis has been constructed in such a way that it consists of a wide variety of the more elementary decision analytic procedures. The purpose of this case study is to provide the reader with a general overview of what decision analysis is all about.

It is written up in the form of a dialogue between the decision maker and a decision analyst; this style of writing has been demonstrated to effectively communicate these ideas.

There are currently plans to use this first draft of the handbook for an intensive course in the Defense Intelligence School and for shorter courses with analysts in the Defense Intelligence Agency and at the Naval Intelligence Support Center. We are attempting to make arrangements for this early version of the handbook to be printed, as a series of separately bound chapters, at the printing office within the Defense Intelligence Agency. Of the currently planned sixteen chapters of the handbook, the introductory case study and six of the most elementary chapters will be printed in a higher volume as it is anticipated that they can be used by students and analysts who desire only an elementary introduction to decision analysis. We anticipate

that these seven chapters will all be written and printing will begin during the next quarter .

2. TASK 1 - RESEARCH ON THE APPLICATION OF DECISION THEORY

2.1 Action Selection (Primary Decisions)

Two case studies involving primary decisions were carried out during the first quarter of the contract period. They are presented briefly in the following sections and will be described in more detail in a subsequent technical report.

2.1.1 Case Study 1. Decision Analysis with Probabilities and Utilities. The first case study considered which of two platforms should be used for the airborne collection of information about a particular country. A previous detailed analysis had shown that a new platform was superior to an older one, currently in use, on all dimensions except political. That is, the new platform cost less, involved less military risk, and was more efficient in gathering information. Because of its greater efficiency, the new platform had been proposed two times previously as a substitute for the old platform. However, on the basis of intuitive judgment of the greater negative political impact of the new platform, i.e., it would constitute a substantial change in the status quo, high-level policy makers had rejected this proposal. Consequently, on the third analysis, a staff man was charged with the task of providing a quantitative evaluation for all components of this decision.

With the aid of DDI, the policy staff man made the following decision analysis of the political implications of using the new versus the old plane. A simple decision diagram of this analysis is shown in Figure 1. Before beginning the analysis, the staff man indicated that he had made previous unsuccessful attempts and stated that he




Figure 1.

did not want to draw a complicated "10,000-branch tree". He wanted to feel that he was driving the analysis rather than that the analysis was driving him due to its great complexity. Consequently, the decision tree was simplified to the simple five-branch tree displayed in Figure 1. The top decision branch represents flying the older plane, the one currently in use, and the lower decision branch represents switching to the new plane. Given that the old plane is being used, an important event which could impact on its political sensitivity is that the country

over which it is being flown could attempt to shoot it down. Consequently, the fork following the decision to fly the older plane consists of two branches. The top one indicates that there is no attempt to shoot the reconnaissance plane down, the lower fork represents that there is such an attempt. Given that such an attempt is made, it can result in a hit or a miss. Since the older plane has been in use over this country for some time, the primary event that would influence its political sensitivity would be a shoot-down. An attempted shoot-down would offer no danger to the new plane because its performance capabilities make it immune to such an attempt, but since its introduction represents a change in policy, it could be interpreted as an aggressive act. Consequently, the fork following the introduction of the new plane consists of one branch which indicates that there is a violent political reaction, and the second branch which indicates there is no political reaction.

Figure 2 is an extension of Figure 1 with judgments of probability and utility appended. For reasons of classification, the numbers representing such judgments are camouflaged but still represent the important idea for purposes of this report. First consider the utility judgments. A value of 0 represents status quo with respect to political impact and a value of -100 represents the worst possible outcome which, in the opinion of the staff man, was the second branch indicating the older plane is shot down. A violent political reaction to the introduction of the new plane represents the second most serious end point and,

in the judgment of the staff man, was approximately 30% as bad as the shoot-down of the old plane. Consequently, it was assigned a




Figure 2.

value of -30. Also, an attempted but unsuccessful shoot-down of the old plane was considered to have only a mild negative political impact, about 10% as bad as a successful shoot-down.

The probability assessments on Figure 2 were supplied by appropriate intelligence analysts. The probability of an attempted shoot-down and the probability of a political reaction were supplied by a political analyst and the probability of a hit given that there is an attempted shoot-down on the old plane was supplied by a technical weapons system analyst.

As shown in Figure 2, it was unlikely that there would be an attempt to shoot down the plane currently in use (20% chance), but

in the case such an attempt was made, it was very likely (90% chance) that it would be successful. Furthermore, it was judged to be very unlikely (10% chance) that there would be any serious political reaction to the introduction of a new plane. The expected utilities indicate the relative political merits associated with the decision to use the old versus the new reconnaissance plane. This analysis shows that, from a political point of view, it is about six times worse (18) to fly the old plane than the new plane (-3).

A sensitivity analysis indicated that from a political point of view, the new plane was less negative than the old plane for considerable changes in judgments of either utility or probability. That is, a change in recommended decision occurred only if the probability of an attempted shoot-down was decreased to below any value that the intelligence analyst would accept or if the negative utility associated with the shoot-down was made less negative in relation to the "reaction" utility than the staff man was willing to accept.

The staff man then presented this analysis to analysts of other agencies in order to verify the validity of the structure of this decision diagram as well as the reasonableness of the assessments of probability and utility. Determining that analysts in other agencies were in general agreement, he presented a briefing to the policy maker in which he used this decision diagram along with his analyses of costs and benefits for the two aircraft.

The decision maker rejected the entire decision analysis concerning

political impact on the grounds that such a complex factor cannot be reduced to quantification. Even though he accepted the remainder of the cost-benefit analysis, he recommended the continued use of the older plane on the grounds that introduction of the new plane might result in a situation that was politically unacceptable. The policy maker, who has the reputation of being exceptionally intelligent, was sufficiently disillusioned by the combination of "wiring-diagram and numbers" that he actually refused to discuss the serious political consequences of a successful shoot-down of the old plane.

The contractors were not present at the briefing to the policy maker but the report of the staff man suggests that successful use of this tool will require some form of education, at least for some of the policy makers. It is interesting to note that the staff man in question is continuing to use this kind of quantification to recommend decisions to the policy maker to whom he reports.

2.1.2 Case Study 2. Decision Analysis with Probabilities Only. The second attempt at a decision analysis was conducted with intelligence analysts instead of with the decision maker. Consequently, the primary input to the analysis consisted of probabilities rather than assessments of value or utility. The problem analyzed was the U.S.'s concern over treaty negotiations with a particular country and whether a treaty would be negotiated successfully during the coming year. A variety of items, including the policy of the United States, will interact on the probability of successful negotiation and

consequently, the analysts constructed three different probability diagrams, each conditional upon a particular U.S. policy. The goal was to display to the policy maker the manner in which the probability of a successful treaty negotiation changed as a function of different U.S. policies. The most aggressive U.S. policy, of course, led to the highest probability of negotiating a treaty, but, because that policy included more concessions, it led to the highest probability of acquiring the least desirable of the possible treaties. This resulted in a very incomplete decision analysis because it is left to the decision maker to intuitively make the trade-off between the probability and desirability of the treaty.

Each probability diagram considered a variety of events not under control of U.S. policy that would impact on whether or not the treaty negotiation was successful. For a complete review of the results of this analysis see CIA Intelligence Memorandum No. 2438/72.

This type of analysis sidestepped a major obstacle to conducting formal decision analyses within DOD. A complete decision analysis requires the assessment of a value structure as well as the assessment of probabilities. While there is little hesitancy on the part of an analyst probabilistically to forecast future events, since that is a central function of his job, there are two serious obstacles to assessing the value structure describing the relative attractiveness of those events to the decision maker. The first obstacle is that values must typically be assessed by high-level policy makers whose time is already overcommitted

and secondly, there are political reasons that argue against permitting a decision maker to expose his value system: If value dimensions relevant to a decision analysis include such factors as money, men, equipment, U.S. prestige, etc., it may be politically uncomfortable for a policy maker to express his trade-offs among those dimensions.

The initial goal was to use the treaty negotiations as a preliminary analysis and then to work with relevant analysts within the National Security Council to conduct a complete decision analysis, i.e., to assess more accurately the values for each of the end points of the probability diagram. Before work was begun with analysts within the National Security Council, events occurred which made it almost impossible that a treaty could be negotiated successfully this year and work on this case study was terminated.

We are currently planning to conduct on-line decision analyses in two other organizations: the Office of Net Threat Assessment within the office of the Assistant Secretary of Defense for Intelligence and the Mideast Task Force in the Office of the International Security Affairs within OSD.

2.2 Information Seeking Decisions (Secondary Decisions)

The type of decision discussed in Section 2.1, "Action Selection", is called a "primary decision". Decisions to acquire information that have the potential of improving a primary decision are called "information decisions". The primary decision is therefore what drives an information decision, but an information decision must occur before the primary

decision is made (but not before it is identified), in order that the acquired information has a possibility to influence the primary decision.

As described in the proposal for this research, the standard decision theoretic analysis of an information decision includes four distinguishable components. The first is the information decision: Should information be acquired, and if so, what kind and how much?; second, the acquisition of the information; third, dependent upon the outcome of the information, the primary decision is made; and finally, some state of nature occurs and, together with the decision made, results in a particular outcome that has value to the decision maker. A general principle is that information is purchased only to the extent that it has a potential for improving the expected value of a primary decision by an amount that is greater than it costs to collect the information.

2.2.1 Difficulties in using the primary-decision approach.

An important class of information decisions within DOD involves collecting information for intelligence agencies: What type and how many platforms should be used, should a new expensive collection system be purchased, etc.? How is it possible to substantially cut the collection budget without doing serious damage to the value of the information being collected? It turns out that the relevant primary decisions are frequently difficult to isolate for these kinds of collection information decisions. One of the problems is that the primary decisions which will be made in the future are not known presently, and secondly, it is frequently


difficult to understand just what impact information will have on a primary decision even if it is known what that decision will be. Another problem is that the man who makes the information decision may not have the information necessary to evaluate the primary decision. This may not be a good state of affairs from a decision theory standpoint, but it is necessary to construct a methodology that will operate within this constraint. Therefore, even though it is useful formally to consider the role of primary decisions for many information decisions, it is important to develop a procedure for measuring the value of information in the absence of explicit consideration of the primary decisions upon which that information will impact.

2.2.2 An approach based on direct assessment of value - use of value diagrams. The words value diagram or tree are used as a name for the theoretical model that has been chosen to measure the value of information in situations where it is inappropriate to consider primary decisions explicitly (this approach has also been called, "goal-dependent utilities" and "goal fabric analysis").

The basic idea of using a value diagram to assess the value of information is to construct a hierarchical structure in such a way that it is possible to measure or assess the relative values of different collection systems with respect to the overall goals of collection. This is done by evaluating the amount of information that each of the collection systems contributes to alternative subgoals and then evaluating the importance of the subgoals with respect to the major goals.

GOALS AND SUBGOALS

INFORMATION COLLECTION



information value

.120
.090
.030
.045
.015
.700
<u>1.000</u>




Figure 3.

Figure 3 is a schematic diagram of an actual analysis that is intended to illustrate how a value diagram works. The object of the study was to assess the relative value of the information being collected by each different, separately fundable collection system or platform, with a particular emphasis on evaluating the air platforms. The ultimate purpose of the study was to reduce the current budget for air collection in such a manner that it would have a minimal decrease in the value of information.

The left-hand side of Figure 3 indicates the goals and subgoals of the information collection and the right-hand side indicates the disciplines and platforms that are responsible for collecting information relevant to the goals and subgoals. The numbers on each branch are importance weights. Within each fork, the importance weights sum to 1.0. With respect to the goals and subgoals, the importance weights indicate the relative importance of different geographic regions, of strategic versus tactical information, and of air versus naval versus ground versus missile information. On the right side, the weightings, again summing to 1.0, indicate the assessed relative values of information being collected by the different disciplines and then by the different systems within each discipline.

Consider first the geographic regions. The weights attached to the branches imply that it is twice as important to obtain information about region A as about region C and four times as important to obtain information about region A as about region D. Thus, these weights

are measures of relative importance of information in different geographical regions. The next branch indicates that, within region B, that information about a strategic threat is four times as important as information about tactical threats. By contrast, within geographic region D, the information about the strategic threat is assessed to be equally as important as information about tactical threats. Similarly, assessments of the relative importance of strategic and tactical information have also been included for geographic regions A and C. However, rather than displaying the complete analysis, the diagram in Figure 3 illustrates how the analysis applies to a single path through the tree.

Continuing with tactical information in region B, information about ground forces is the most important, followed by air, and finally naval and missiles, which are equally important.

The collection part of the tree is divided into disciplines on the left and platforms on the right. With respect to disciplines, Figure 3 shows that the value of photographic intelligence currently being collected against naval tactical forces in geographic region B, is eight times as great as the value of information being collected about communications, and ten times as important as information being collected about radar activity in that area. Finally, information provided by attachés is assessed to be the least important in this area. Next, consider the final column; 90% of the value of photographic information is currently being obtained by airborne collection platforms, and 10% by nonairborne systems. Of the airborne platforms, airplane A contributes 50% of

the value and the combination of airplanes D and E contributes the other half. For information about communications, on the other hand, only 40% comes from airborne platforms, and 70% of the value of this information is being obtained from airplane B. Notice that only airplane A contributes both photographic and communications information whereas B and C contribute only communication information and planes D and E contribute only photographic information.

The small tree on the bottom of Figure 3 shows the relative values of information collected by the different systems and platforms as implied by the larger value tree. The relative value of a platform is a weighted sum that is equal to the value of the information contributed to a goal or a subgoal weighted by the relative importance of those goals, and then added across all goals. For example, the relative value of the photographic information that airplane A collects about tactical naval forces in geographic region B is equal to the product of all the weights on that path ($.30 \times .20 \times .10 \times .40 \times .25 = .00006$), and so on. The total relative value of information being provided by airplane A, as displayed in the bottom tree on Figure 3, is equal to the sum of all of its individual weighted values (totaled to .12), only two of which are shown in Figure 3 (.00216 and .00006). In this case, the column titled "information value" shows that airplane A contributes 12% of all information being collected by all collection systems in all geographic regions.

2.2.3 Value diagram - first case study. We have applied

the value tree as a case study in DIA for the purpose of measuring the relative values of information being produced by air platforms in one particular geographic region. The structure which developed from this case study is far more elaborate, but analogous to the one displayed in Figure 3. Different analysts were used to structure the value tree and to assess the relative importance weights for different branches on the tree. For example, one set of analysts assessed the relative importance of subgoals such as strategic and tactical forces. The air analysts assessed the relative contribution of the different collection disciplines for air, whereas naval analysts assessed the relative importance of the different disciplines to their problems, etc.

Problems emerged when these intelligence analysts were faced with the task of assessing the relative contributions of different platforms to the collection disciplines. The reason is the following: the air analyst, for example, was aware of which collection discipline contributed to the value of his information but typically would not know what platform provided that information. In order for him to make this assessment, it was necessary for the collection analyst to inform the intelligence analysts which airborne platforms, for example, were collecting over which airfields. The intelligence analyst was aware of the relative value of the information that he was receiving concerning each airfield. Therefore, when he was able to pair up specific intelligence targets with each of the collection platforms, he was able to assess

the relative value of the information currently being derived from each of the collection platforms.

Thus it was necessary in these cases to bring together the analyst who knew the value of information being collected on certain targets with the analyst who knew which platforms were being flown against those targets. This was done in order to assess the relative value of the information derived from the different platforms.

This analysis is now being used to search for "soft spots" in the allocation of airborne reconnaissance resources. That is, the relative information values of airborne platforms are being compared with the corresponding per cent of dollar costs of those platforms. While the correspondence should not be perfect, low information values in per cents together with large dollar costs in per cents will be used as a flag for possible changes or actual elimination of certain platform usage.

Notice that this is an open-form analysis as was discussed in the proposal for this research. This analysis in no way prescribes to the decision maker explicitly how he should allocate his resources, but it does provide him with ready access to information that is useful in making that decision.

There are several ways of extending this methodology for the assessment of the relative value of information. One method is to address the information decision problem directly. Suppose, for example, that the collection officer has the option of either using airplane A,

substituting airplane A', or completely eliminating that platform.

One way of evaluating the alternatives is to rerun the relevant portions of the analysis for each of the three alternatives; platform A, platform A', and no platform.

For such an analysis, it is necessary to use a common standard across all platforms. One possibility for that standard is the current value of information. Reconsider Figure 3. Assume that for region B and tactical naval information platform A' is assessed as being capable of providing 80% of the current photo information, and twice as much of the current communications information as was platform A. The weighted values of platform A' are then calculated by multiplying the weighted value of photo for platform A by .8 and multiplying the weighted value of the communications information for platform A by 2.0. It is then possible to derive an assessment of the change in value of information as well as an assessment of the change in cost of information for each of the alternative information decisions.

During the next quarter, we plan to continue this development of a theoretical methodology to be used for information decisions, and to test this methodology by applying it to resource allocation decisions within DIA.

2.2.4 Value as a function of decision variables - second case study. A second procedure for evaluating information decisions in the absence of a primary decision has also been developed, and evaluated in a case study form. This procedure involves an assessment

of the manner in which both value of information and cost of information change as a function of the number of hours or number of missions for a single airborne platform. This problem developed from a case in which an airborne platform had initially flown missions for a total of 80 hours per month, and later flown missions at a rate of 30 hours per month. The drop in information value commensurate with the drop in hours per month was greater than had been anticipated. The task was to predict from the available data the value and cost of information at different levels of activity, such as 60 hours per month. This can be viewed as a subordinate methodology to that previously described where the goal was to determine the relative value of information derived by changing platforms.

In this particular case, it was relatively straightforward to measure the reduction in cost resulting from the reduction in flight hours. It was, however, anticipated that the value of information would increase in a negatively accelerated manner as a function of increase in flight hours. That is, the first few flight hours would result in a rather large increase in the value of information, whereas additional hours would result in smaller increments of value. Instead of an increase, the intelligence and collection analysts assessed that the value of information decreased by 75% with a reduction from 80 to 30 hours, whereas the cost of flying the missions decreased by only about 15%.

It turned out that the reason for the unexpected sharp decrease

in the value of information was that there were far more unidentified intercepts with the 30-hour month than with the 80-hour month. That is, it was necessary to build a data base up to some certain level before the additional data being collected would be of any value. This result led to the speculation that, in the case of communications intercepts, an S-shaped function generally could be expected when assessing the relation between value of information and number of missions, whereas a negatively-accelerated function would relate cost to number of missions. This suggests that there will be a certain region of activity where the value of information changes rapidly for small increments in cost. The optimal level of activity, of course, will depend upon available resources as well as alternative collection systems which can be used. The implication for a general methodology for information decisions in intelligence is that it is critical to look for non-linearity in functions relating information value as well as cost to level of activity of each collection system.

2.3 Supporting Technical Experiments

Two experiments on probability assessment were conducted during the first quarter at the Defense Intelligence School using approximately one hundred student analysts as subjects. Each experiment required approximately an hour-and-a-half and each subject served in only one of the two experiments.

2.3.1. Experiment 1 - Odds versus probabilities as response mode for assessing the likelihood of categorical events. This experiment was not included in the original proposal but its subsequent

inclusion was discussed with and approved by the scientific officer.

Previous psychological research has shown that when a person is asked to estimate the likelihood that a statement or an event is true, the value that he assesses depends upon the response mode being used. For example, the individual who assesses that there is an 80% chance that some event will occur, may assess that it is ten times more likely the event will occur than it will not occur. These two assessments are, of course, inconsistent, (i.e., .80: .20 or 4:1 vs. 10:1), but it is not obvious which of the two modes, probabilities or odds, yields a better measure of the state of knowledge of the probability assessor.

Previous research that has compared probabilities with odds was conducted in an environment where the task was to revise likelihood assessments in the light of new information. A general finding of that research has been that a subject's estimates tend to be lower than they should be and that the odds are more extreme than the probabilities. Therefore, the odds estimates tend to be closer to optimal than are the probability estimates.

The problem is that this observed superiority of odds over probabilities may be artifactual. If it is possible, for example, that in a static situation where it is not necessary to update assessments in the light of new information, assessments are not conservative or low. In this situation, odds, if they remain more extreme than probabilities, may provide overestimates of the subjects' state of knowledge. Consequently, Experiment 1 was designed to evaluate, in a static

situation, whether odds are more extreme than probabilities and also whether they are more accurate.

Stimulus material consisted of fifty factual questions such as, "There is a higher per capita income in either: (1) Washington, D.C., or (2) The state of California." The subject's task was to indicate which of the two answers was correct and his confidence in that answer. His confidence assessment took the form either of probabilities or odds.

Forty subjects were randomly divided into two groups. In Group 1, the subjects completed all fifty questions first with probability estimates and then went through them with odds estimates. In Group 2, the subjects first used odds estimates and then used probability estimates.

Preliminary analysis of the data has begun and we plan to conduct three analyses. The first will compare the magnitude of odds estimates and the corresponding probability estimates for each subject and each question. Previous research suggests that odds will tend to be more extreme. The second analysis will use a scoring rule to compare the relative quality of each estimate. This will be possible because the experimenter knows which of the two answers, such as, "Washington, D.C." or "The State of California," is true.

The final analysis will evaluate the calibration of odds estimates with the calibration of probability estimates. This will be done by ranking all probability estimates in order of increasing magnitude across all subjects and questions and then assigning those estimates

to one of five categories determined by the magnitude. Within each category, the average estimate will be compared with the per cent of correct answers. Perfect calibration implies that the average estimate should equal the corresponded per cent correct and if subjects tend to be conservative, the average estimate should be smaller than the corresponding per cent correct. If subjects are excessive, their average estimates should be more extreme than the corresponding per cent correct. This measure of calibration will be performed separately for probabilities and for odds, and will provide a second means for comparing the relative quality of the two as response modes. If the results are sufficiently interesting to warrant it, the experiment will be written up and prepared as a technical report.

2.3.2 Experiment 2 - Probabilities of points versus intervals along a continuum. This second experiment was described in detail as Experiment A of the original proposal. Its purpose was to evaluate the relative merits of two different response modes that could be used for assessing the probability of events that lie along a continuum. Examples of such continua are the cost of a proposed weapon system, the speed of an approaching airplane, the number of men that will be lost in a planned mission, range, speed, altitude and accuracy of a particular missile, and the magnitude of the military portion of the R & D budget for a particular country.

Various assessment procedures have been devised and fall into two general classes. One involves the assessment of the relative

probabilities of points along the continuum and naturally yields a continuous probability distribution or a density function. The second involves the assessment of the relative probabilities of intervals. Since it is possible to derive a cumulative distribution from a density function, or a density function from a cumulative distribution, the class to which a procedure belongs does not necessarily determine which of the procedures should be used in any particular instance. However, most business texts recommend the assessment of fractiles which are then converted to cumulative probability distributions.

As an example of the fractile procedure, a probability assessor may be asked to assess the median of his probability distribution, such that he feels the true answer is just as likely to fall above as below that value. He then assesses the range such that he is 90% sure that the true answer will fall within and 10% sure that it will fall outside.

A more direct procedure for assessing a probability density function involves the assessment of the relative probabilities of particular points along a continuum. For example, the probability assessor could select the mode as the most likely event along the continuum and then select points above and below that mode which are one-half or one-fourth or one-tenth as likely as the mode. It is then possible to draw a curve through those points to represent the density function.

Stimuli for this experiment consisted of such questions as, "How long has Tito been the leader of Yugoslavia?", "What is the number of total operational B-52's currently in the U.S. inventory?", "What

is the FY1973 U.S. DOD budget?", and, "What is the number of men in the U.S. Army on active duty?" Notice that the answer to each of these questions lies along a continuum and, in each case, the correct answer is available. Each subject used two different procedures, one focusing on intervals and the other on points, to answer each of the questions. With the interval procedure the subject trisected the continuum. That is, he divided the continuum, such as number of men, into three intervals such that he expected that it was equally likely that the true answer would fall in each interval. For the point procedure, the subject first selected the mode, or the most likely answer, and then selected the point above and point below the mode that was half as likely to be the correct answer. Seventy subjects participated in this experiment.

The data analysis will consist of a comparison of the two responses, points versus intervals, with respect to quality of calibration. This will be done by first identifying a high, a medium, and a low interval for each response.

For the interval procedure, the interval boundaries will be defined by the trisection, i.e., the three intervals into which the subject divides the continuum. For the point procedure, the middle interval will be defined as that interval between the points above and below the mode that are half as likely as the mode. For the interval procedure, of course, the subjects will be well-calibrated only if one-third of the correct answers fall into each of the three intervals.

For the point procedure, it turns out that approximately 75% of the area of the density function falls between the points above and below the mode that are half as likely as the mode.

A preliminary analysis indicates that, for the interval procedure, about one-fourth of the correct answers fall in the middle interval, whereas for the point procedure, about 40% of the points fall in the middle. Thus, too few of the correct answers lie in the middle interval for both procedures. Generally, the subjects have made the center interval too small, indicating they are more sure of the correct answer than their knowledge warrants. However, the degree of bias is considerably greater for the point than for the interval procedure.

Since this first analysis favored the interval procedure, a second analysis provided a second evaluation of that procedure. Consider a set of credible intervals that vary in width. It should turn out to be the case that the average error or distance from the estimate of central tendency, such as the mode, to the true answer, increases as the size of the credible interval increases. That is, the probability assessor should assess a wider credible interval when he knows less about the true answer and therefore his best estimate turns out, on the average, to be in greater error.

In order to convert all questions to a common unit, credible intervals across subjects were converted to standard scores within each question and the amount of error was converted to standard scores within each question. Each pair was treated as a unit and sorted according

to size of absolute error in standard scores and then grouped into categories beginning with the highest 10% of the questions, the next-highest 10%, etc. Then the mean absolute error, and the mean size of the credible interval associated with that error were calculated for each of the ten categories. The resulting correlation was .93, indicating that about 87% (.93 squared) of the variance in the credible interval width can be accounted for by the mean error of the mode. As should be the case, subjects estimated wider credible intervals when they were more in error about the true answer.

The results of this experiment are now being written up as a technical report and should be available by the end of April.

2.3.3 Experiment 3 - Decomposition of the assessment of continuous variables. Subjects using the interval procedure in Experiment 2 evidenced sufficiently good calibration, eliminating the need to conduct Experiment B of the original proposal. The intention of that experiment was to evaluate procedures for training probability assessors with respect to calibration only if subjects appeared to be poorly calibrated with the interval procedure. Consequently, we plan to conduct the experiment identified as Experiment C in the proposal as the third experiment under this contract, and will perform it during the summer of 1973. That experiment is designed to evaluate procedures for breaking the assessment of credible intervals into component subtasks.

3. TASK 2 - WORKSHOPS

Task 2 of the contract was to conduct workshops on decision analysis with on-line DOD decision makers participating. A potential list of participants has been identified and we plan to conduct those workshops during the Fall of 1973, using the results of research conducted during the year.

4. TASK 3 - HANDBOOK FOR USERS OF DECISION ANALYSIS

The proposed handbook for decision analysis now consists of the sixteen chapters in varying degrees of completion as shown in Table 1.

The greatest amount of effort has been devoted to the first chapter: a case study intended to provide an overview of decision-analytic procedures. This study features a naval engagement in which a task force commander must decide whether or not to free his weapons to shoot an approaching unidentified airplane. It includes structuring a decision diagram, solving that diagram, revising probabilities in the light of new information, the use of probability diagrams to assess probabilities, the use of multi-attribute utilities to combine several different dimensions of value, and how to assess the value of an information source. This chapter will be refined and printed during the second quarter.

This case study, together with the other six chapters identified by asterisks in Table 1, are planned for use in courses on probability analysis to be conducted within the Defense Intelligence Agency and the Naval Intelligence Support Center. Therefore these six chapters will also be refined and printed during the second quarter.

Current arrangements are for all sixteen chapters to be printed

by the Defense Intelligence Agency. It is because of the combined need for these chapters in the Defense Intelligence School and in DIA that the DIA print shop has tentatively agreed to print the handbook chapters at no cost to the present contract or its follow-on. The chapters will be printed as individual packages over the next several months with each chapter requiring three weeks for printing after it is received.

Table 1

Handbook for Decision Analysis

- * Case Study - an example of the use of the tools of personalistic decision analysis
- Structure of a Decision Tree - setting up the elements of a decision problem for solution
- Solving a Decision Tree - averaging out and folding back
- Combining Dimensions of Value - outcomes which have multiple criteria
- Attitude Toward Risk: Utility Analysis
- * Personal Probabilities - the meaning of probability and the rules of probability theory
- * Direct Assessment of Categorical Probabilities
- * A Scoring Rule for Probability Assessment - a criterion for a good probability assessor and the scoring rule test
- * Probability Distribution for a Continuum - methods of assessing uncertainty about continuous quantities
- * Probability Diagrams - decomposing the diagnosis problem
- Pruning Probability Diagrams
- * Inference from Evidence: Bayes' Theorem - the impact of evidence on the likelihood of hypotheses
- Bayes' Theorem and Continuous Distributions - special techniques for assessing likelihood ratios

raw material	rough copy	final copy	printed copy
X	X		
X	X		
X			
X			
X			
X	X		
X	X		
X	X	X	X
X	X	X	X
X	X		
X	X		
X	X		

* Planned for use in DIA and NISC courses

Table 1 (continued)

Hierarchical Inference - the impact of data on
the likelihood of hypotheses when
the intervening events are uncertain

Information Decisions: Impact on Primary
Decisions - the correct price to pay
for information given the possible
effect on the decision to be made

Information Decisions: Value Tree

raw material	rough copy	final copy	printed copy
X			